



Information aggregation in a prediction market for climate outcomes

Elmira Aliakbari ^a, Ross McKittrick ^{b,*}

^a Fraser Institute, Vancouver, BC, Canada

^b Department of Economics, University of Guelph, Guelph, ON, Canada

ARTICLE INFO

Article history:

Received 16 November 2017

Received in revised form 13 February 2018

Accepted 1 June 2018

Available online 5 June 2018

JEL classification:

Q54

Q58

H23

G13

Keywords:

Climate change

Uncertainty

Carbon tax

Tradable permits

State-contingent pricing

Prediction markets

ABSTRACT

Two forms of uncertainty in climate policy are the wide range of estimated marginal costs and uncertainty over credibility of rival information sources. We show how a recently-proposed solution to the first problem also helps address the second. The policy is an emissions tax tied to average temperatures, coupled with permits that exempt the emitter from paying the tax in a future year. It has been shown that the resulting tax path will be correlated with future marginal damages. It has been conjectured that the permit prices will yield unbiased forecasts of the climate, which, if true, would address the second uncertainty. We confirm the conjecture by describing a trading mechanism that converges on unbiased forecasts if traders are risk-neutral. Risk aversion slows down but does not prevent convergence. We also show that the forecasts are more likely to be sufficient statistics the stronger the consensus on climate science.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

A policymaker wanting to address the climate change issue faces many information problems. This paper examines novel solutions to two of the most acute: uncertainty regarding marginal damages of greenhouse gases, and incomplete or biased selection of information for presentation to policymakers. The first derives in part from scientific uncertainty over Equilibrium Climate Sensitivity (ECS), which is defined as the projected long term climate response to doubling the concentration of CO₂ in the atmosphere, after enough time has passed for all components of the climate system to adjust. A widely-used study by Roe and Baker (2007) yielded an ECS distribution with a 5th percentile of 1.72 °C and a 95th percentile of 7.14 °C (IWG 2010). Applying this distribution and a 3% discount rate, Integrated Assessment Models (IAMs) generate a range of marginal damage estimates spanning -\$22 to \$727 per tonne of CO₂ (IWG, 2013). Hence, based on mainstream climate and economic modelling, the optimal climate policy seems to lie somewhere between a small subsidy for, and an effective ban on, all emissions—a rather unhelpful state of affairs to say the least.

Past attempts to reduce the uncertainty of climate policies are reviewed in Section 2. The approach considered herein involves a

temperature-indexed carbon tax (McKittrick, 2011) that presents emitters with a dynamic pricing rule rather than a commitment to a specific price path. On an ex-post basis, the rule yields a price path highly correlated with the unobservable true marginal damages trajectory, thus contributing to resolution of the first form of uncertainty, since firms subject to the tax must form unbiased expectations about future values and plan accordingly. A proposed modification to the policy instrument can also address the second information problem. To facilitate expectations-formation, Hsu (2011) proposed pairing the tax with a sequential futures market for tradable certificates, each of which would exempt the holder from paying the tax on a tonne of emissions in a specified future year. He conjectured that since firms have a financial incentive to get the forecasts right, the price path thus generated would provide the most objective and informative forecast of future marginal damages and, by implication, future climate warming, since investors will have an incentive to use all available information, i.e. to avoid cherry-picking.

This paper is concerned with the conditions under which the permits trading mechanism would yield informative and unbiased forecasts of future climate states. We refer to the temperature-indexed pricing rule as a state-contingent carbon tax, and the trading system as a futures market for exemption permits. Since projected damages are expected to occur rather far in the future, they cannot be directly measured and will depend strongly on the modeler's choice of ECS,

* Corresponding author.

E-mail address: ross.mckittrick@uoguelph.ca (R. McKittrick).

among other parameters. Policymakers must therefore rely on damage estimates that rest on expert opinion, which raises the possibility of bias since scientists will be aware that their parameter selections have implications for policy outcomes over which they may have preferences (Johnston, 2012, InterAcademy Council, 2010). A prediction market that generates a state-contingent carbon tax path based on expectations among agents who have a financial incentive to get their forecasts right would potentially resolve this problem by creating an incentive to form unbiased climate projections. An ideal auction price should therefore be both an unbiased estimate of the actual future price (and hence the actual future climate state), and also a sufficient statistic (making efficient use of all available information).

We model a prediction market implementing the McKittrick (2011) and Hsu (2011) state-contingent carbon tax/prediction market system in an auction framework developed by Kyle (1985) and Foster and Viswanathan (1996, herein FV96). Three types of traders are assumed to participate. Risk-neutral firms subject to the emissions tax are assumed to have private access to noisy signals about the likely future climate state and hence the likely future tax rate, and to make bids for permits based on their profit-maximizing strategies. Also, an unspecified number of uninformed traders generate noise in the market by trading based on purely random signals about the future climate state rather than informative private signals. A market maker who only sees the aggregate order flow but does not observe individual bids or the number of uninformed traders clears the market in each round of bidding, thereby generating a price signal which is incorporated into information sets by traders in subsequent trading rounds.

Interestingly, this trading set-up corresponds to that of a newly-proposed climate prediction market to be run by a UK-based financial firm (Winton Capital, 2017). Traders will be able to buy and sell securities whose underlying value will be tied to future values of publicly-available global temperature measures. Though there is no state-contingent policy needing to be hedged, the market designers were motivated by concerns that biased information interfered with optimal policy planning, asking: “Are peoples’ positions driven by what they want to be true or what they believe to be true?” (Winton Capital, 2017) The proposed market allow informed and uninformed people to trade certificates tied to future climate states, with Winton acting as a non-profit market maker.

Our paper provides some evidence that such a market will in fact aggregate information in a useful way. Employing a proposition due to FV96 we confirm part of the conjecture of Hsu (2011) for the case of risk-neutral traders: such a prediction market would yield an unbiased forecast of the future tax rate, and hence the state of the climate. It is less clear that the prediction market will use all available information efficiently. Interestingly, one’s belief about the level of consensus around climate science strongly influences one’s interpretation of the market outcome. The question of the degree of consensus in climate science, and what specific points experts actually agree on, is itself controversial (Berry et al., 2016; Tol, 2014). We treat the notion of scientific consensus as corresponding to a high level of correlation of informed traders’ signals, and hence of their beliefs. In other words, a strong scientific consensus means that regardless of the source an informed trader draws upon, it is the same information as other traders are receiving. We find that, even in the presence of noisy uninformed traders, the market price will converge to a sufficient statistic (in other words, an expectation that efficiently uses all available information) if, at the start of trading, private beliefs of informed traders are highly correlated, or in other words, if there is a truly strong consensus on the scientific issues, so that traders seeking credible information are effectively drawing the same signals regardless of source. In this case, the trading process we describe will yield a price path that incorporates all relevant information and no trader would be able to improve on the market price forecast using his or her private information set. Conversely, to argue that the prediction market

outcome fails to take account of some relevant information about the future path of the climate requires an assumption that there is no current consensus about climate science.

The FV96 framework assumes traders are risk-neutral. We extend the theory by allowing firms to be risk-averse. We show that risk-aversion slows down the convergence process, so the market outcome is more susceptible to the influence of uninformed traders.

Prediction markets were first proposed in 1988 by researchers at the University of Iowa interested in predicting US presidential outcomes (Segol, 2012). The markets were designed based on the informational role of prices and the efficient market hypothesis suggested by Hayek (1945) and Fama (1970). They are not magic devices for generating error-free information about the future, since they are still susceptible to all the biases, wishful thinking and other forms of irrationality that enter human judgment. Their advantage is that they create a financial incentive for participants to be objective and to use all available information, and this leads us to ask how they might be harnessed to yield unbiased and efficient climate forecasts. Although market efficiency assumptions have been exposed to some critiques, especially from behavioral economists, other researchers (Wolfers and Zitzewitz, 2004; Berg et al., 2008) have shown that prediction markets perform well and generate accurate forecasts. A recent survey by Graefe (2016) compares prediction markets to a variety of polling and prediction methods in international elections and shows that they consistently provide the most accurate forecasts (even though, as noted in Graefe, 2017, they can still fail, as they did in the Brexit and US 2016 Presidential election cases). This motivates our interest in using them for predicting climate change, which is an application of considerable international importance.

Recent work by Atakan and Ekmekci (2014) has shown that a common-value market may fail to aggregate information if the buyers must simultaneously choose an action, the bid payoff is jointly determined by the action chosen and the state of the world, and the price of the object is independent of the state of the world. In our set-up buyers do not commit to any action other than holding the permit, the value of which is fully determined by the state of the world independent of any subsequent action chosen by a bidder. Consequently, the failure of information aggregation in Atakan and Ekmekci (2014) does not occur in our model. However, if the policy game herein were modified such that every bidder must simultaneously commit to an action to be taken after each auction, the value of which would depend on the state of the world after the price is revealed, then we would need to take into account the possibility that an equilibrium may emerge in which no information is aggregated.

The assumed damage function underlying our model is continuous (see Assumption 4 in McKittrick, 2011) as is typically assumed in climate policy models. Lemoine and Traeger (2012, 2016) have explored a model with one or more potential discontinuities in the state function, reflecting concerns about possible “tipping points” or bifurcations in the climate system at which an irreversible transition to a new state function takes place. Such tipping points are associated with changes in the pollutant stock (concentration), which is a heavily smoothed function of annual emissions in the case of carbon dioxide. The possible existence of a future tipping point, the timing until it is encountered, and the magnitude of the discontinuity it might impose, are additional forms of uncertainty that we do not address herein. The prediction market we propose may help reduce these uncertainties since traders will profit from correctly anticipating any form of change in the state variable, but establishing this formally is beyond the current limits of our analysis.

The paper is organized as follows. Section 2 provides the literature review on this topic. Section 3 presents the model, assumptions, and definitions. Section 4 describes the linear equilibrium of the model and its necessary and sufficient conditions. Section 5 presents the results. Section 6 extends the existing model to include risk aversion and the final section briefly summarizes the conclusions.

2. The uncertainty problem

There is a large literature on pricing carbon dioxide (CO₂) emissions as a way of addressing the global warming externality. Using classical theory, the price should be equal to marginal social damages from emissions. However, in the context of global warming, there is significant uncertainty regarding the magnitude of the effect of greenhouse gas emissions on global warming and the magnitude of the impact that global warming has on the economy. Hence identifying the optimal policy has proved to be challenging.

Nordhaus (1991, 1993a, 1993b) introduced the Dynamic Integrated Climate and Economics (DICE) model to help study the mechanisms connecting CO₂ emissions, atmospheric concentration, changes in temperature, and changes in output and consumption, and the effects over time of these things on welfare. Using a growth model he found an optimal path for emission abatement and the associated carbon tax conditional on assumptions about key parameter values. Ever since, many DICE-style Integrated Assessment Models (IAMs) have been developed. The advantage of the IAM approach is that it yields a specific price path. A major disadvantage is the assumption that all the key parameter values are known at the start, and no new information will be obtained over time, so there is no role for learning. The option of learning would imply that a fixed path is not prescribed at the start, only a starting value plus a rule for assimilating new information.

Kelly and Kolstad (1999) relaxed the no-learning assumption by incorporating Bayesian updating about the relationship between temperature changes and greenhouse gas levels in the DICE model. Their main result was that it would take between 90 and 160 years for the uncertainty to resolve sufficiently to reject an incorrect policy choice. Leach (2007) added uncertainty over the time lag of temperature change and found it could take hundreds or thousands of years until uncertainty is resolved sufficiently to reject an incorrect policy stance.

McKittrick (2011) proposed an alternative approach that uses the information implicit in the climate state function by tying the emissions tax linearly to an observed measure of atmospheric temperatures. In this way prices are continually updated, so emitters must form expectations about the future path of warming (and the emissions tax) rather than being given a deterministic schedule by policymakers. If global warming is expected to be rapid, taxes will be expected to go up rapidly, and vice versa. Agents would need to acquire credible forecasts of climate outcomes in order to make long-term investments, and would have an incentive to identify the most accurate information.

In order to provide objective information and allow for hedging of investor risk, Hsu (2011) proposed the idea of coupling the state-contingent pricing rule with a market for permits, each of which exempts the holder from paying the tax on a unit of emissions in a future period. The market for future exemption certificates would motivate buyers to seek information about the likely future climate state, and the trading process would aggregate the information into an equilibrium price. This would provide the best possible forecasts of future tax rates, as long as the trading process yields objective, unbiased price expectations. The price path, in turn, would yield unbiased forecasts of the future temperature path. Our model herein focuses on the question of whether such a mechanism would yield unbiased forecasts and make efficient use of all information available to traders.

3. The model

Following McKittrick (2011), at a future date T , a state variable s_T , which is a non-manipulable measure of the climate state,¹ will be revealed. Based on this observation, the policy maker will impose a tax

rate $\tau_T(s_T)$ which is unique, so $s_T \implies \tau_T$ and vice versa. Following Hsu (2011), instead of paying the tax in year T , firms have the option to submit emission allowances that they already purchased in previous years, each of which allows a one tonne exemption from paying the tax. We assume allowances (which we also call permits) dated for one vintage T cannot be banked for use in another year. Therefore, the spot price of a permit p_T in future year T must equal the tax rate, so $p_T = \tau_T(s_T)$.² Since the functional form of τ_T is known, the permit price also reveals s_T .

At current time t consider a market for tradable permits for vintage year T . We assume there are three types of traders in this market who will trade a specific permit over $h = 1, 2, \dots, k$ auction rounds. First, there are M risk-neutral traders who possess disparate private information about the true climatic state in period T , and thus about the true price of a permit. Each trader's initial private information is denoted $g_{i0}(T)$ where $i = 1, 2, \dots, M$. The private signal can consist both of publicly-available information and privately-acquired information. Second, there are several uninformed traders who do not possess any private information and trade as a single entity in this market. This means that a single aggregate bid emerges from this group. The assumption that uninformed traders are present allows the results to be robust to the possible presence of agents who act on signals that have no valid informational content. As we will discuss below, the assumption is also necessary to avoid the Grossman and Stiglitz (1980) paradox that, if only informed traders are present in a market, information will not actually be collected and used. Third, there is a set of risk-neutral market makers who can be considered as large private financial institutions whose role is to clear the market at zero expected profit. The market makers are responsible for facilitating trades in this market.

At the beginning of the first auction round ($h = 1$), M informed traders receive different noisy signals regarding the true state variable in a future date T . Using their private signals they each form different expectations \hat{p}_{iT} about the future price. Then, informed traders and uninformed traders simultaneously choose the number of permits they want to buy and submit their quantities (market orders) to the market makers. We denote the informed traders' market orders and the uninformed traders' aggregate market order as $q_{ih}(\hat{p}_{iT})$ and u_{Th} , respectively. Since we are only considering a single vintage period T , wherever the T subscript is not needed we will suppress it for clarity. At this stage, when choosing their permit quantity, the only information informed traders have is their initial signal, so their information set is $\Omega_{ih}(g_{i0})$. Market makers observe the aggregate quantities ($y_h = \sum_{i=1}^M q_{ih} + u_{Th}$) without seeing individual orders separately, and they set a price determined by a competitive process that yields zero expected profits for the market makers, implying that it becomes equal to $E[p_T|y_h]$.³ Everyone observes the price and the total quantities traded. Next, market makers trade a quantity $m_h = -y_h$.

In the subsequent auctions, informed traders will trade based on their initial signals and what they learn by observing previous rounds.⁴ In other words, their information set consists of their initial signals, past prices, past market orders, and past quantities traded by themselves, which we write as $\Omega_i(g_{i0}, p_{-h}, y_{-h}, q_{i,-h})$ where $-h$ denotes all previous rounds. Market makers do not observe the individual quantities; when they receive the aggregate quantities, they update their beliefs regarding their estimate of p_T and set the price accordingly. Hence, their information set consists of the current and past aggregate order flows and also past prices, $\Omega_{MM}(y_h, y_{-h}, p_{-h})$. The process of trading goes for

² Since τ_T is a linear and known function of s_T , trading in contracts for the future carbon tax would yield the same finding as trading in contracts for the future global temperature.

³ The market maker's expected profit is $E(\pi) = -y_h (E(p_T) - p_k)$ where k refers to the last round of the auctions. Therefore at each action round h where $h \neq k$, a price equal to the expected value of a permit given the observed order flow, $p_h = E(p_T|y_h)$, would yield zero expected profits for market makers.

⁴ Buying permits on a futures market helps investors hedge against costs of future climate policy.

¹ McKittrick proposes the mean temperature of tropical troposphere as measured by weather satellites as the state variable. Each system for measuring the state variable has its strengths and weaknesses but we set that issue aside for the present study. See Cederlöf et al. (2016) for a comparative survey of climate data sets and their relative uses.

k rounds where k is some finite, large number. At this point when considering the trading process for the vintage year T , we ignore information spillovers that might exist from climatic prediction markets for other vintages. However, this information can be considered part of the publicly-available information used by traders and therefore would not have any effect on our model.

The trading set up corresponds to that applied in well-known financial markets such as the New York Stock Exchange (NYSE) and Tokyo Stock Exchange. The closest application of this set up is the NYSE call auction opening where market makers indicate a price range, traders submit their market orders, and at last market makers adjust the initial prices. In our model we have taken into account strategic behavior of the informed traders given their private information, meaning that they are not price takers and they consider the effects of their actions on the price and informational content of the price when choosing their actions. Since we have a finite number of traders, it is not reasonable to assume price taking behavior. For instance, if information is costly to obtain, there might be some large firms that can invest sufficient amounts to acquire exclusive information, leading them not to behave as pure price takers.

Choosing the number of permits to make available for a vintage year T is a critical matter. The number of available permits should be large enough to create a real market so that traders would have enough incentive to obtain information and engage in trading. However, it cannot be so large as to exceed the emission level associated with the tax that period, otherwise the price would go below the tax rate and converge on zero, defeating the purpose of the hybrid system (Hsu, 2011). In order to avoid this, a policymaker would first need to have an estimate of the emission level in a future year T , e_T , and then set the number of available permits to be a fraction of that.

Inclusion of an appropriate discount rate for future profits will not change the results in our model. In the case of a long-term prediction market, we expect traders to forecast and discount their future tax liability, so that the price of a permit reveals the discounted expected liability.

Following FV96, we assume that the price of a permit for future year T is a random variable with a normal distribution with known mean $\rho(T)$ and variance $\Sigma(T)$, so $p_T \sim N(\rho, \Sigma)$. The “true” price of a permit is a realization of the random variable p_T that is fixed and unknown to all the traders prior to T , however, informed traders observe a noisy signal regarding this true value.⁵

We assume that the signal vector $g = [g_{10}, g_{20}, \dots, g_{M0}]$ is drawn from a multivariate normal distribution with variance-covariance matrix Ψ_0 and zero means. The form of all distributions, including p_T , is known to all traders in this prediction market. We assume that all signals have the same initial variance Λ_0 , the initial covariance Ω_0 between any two signals is the same, and the cross covariance with the true price of the permit c_0 is the same for all signals. As shown in the Appendix, these three assumptions imply:

$$E(p_T | g_{10}, g_{20}, \dots, g_{M0}) = \rho + \theta \bar{g} \quad (1)$$

where \bar{g} is the average signal and θ is a constant. This means that a constant multiple of the average signal, $\theta \bar{g}$, is a sufficient statistic for all the information known by traders (all the signals) to predict p_T . In other words, $\theta \bar{g}$ uses all the available information (all the signals), and

⁵ The distribution of possible future temperatures may not be Gaussian. Roe and Baker (2007) point out that in the standard representation of climate sensitivity, if the feedback parameter is normally distributed the equilibrium climate sensitivity (ECS) could have fat tails. Weitzman (2009) used this result to argue that an insurance premium against catastrophe approaches infinite present value. But Roe and Bauman (2013) criticized Weitzman's argument because it ignores the time-to-equilibrium, which goes up with the square of climate sensitivity. The realized warming associated with the fat upper tail of ECS would take a thousand years to materialize (Roe and Bauman, 2013, p. 653). Consequently, even under a high ECS, the distribution of temperatures on the time scale of the market we consider would not necessarily depart from a Gaussian profile.

knowing any other functions of $g_{10}, g_{20}, \dots, g_{M0}$ will not improve the estimate of p_T . As mentioned before, the normal distribution of the random variable p_T with mean ρ and variance Σ is known to all participants. Therefore, all that is needed to predict p_T is $\theta \bar{g}$. In fact, knowing $\theta \bar{g}$ would provide us with the best estimate of the true value of permit p_T as it is a sufficient statistic for all the available information. This assumption simplifies the complex model and will be used in the following discussion.

It is helpful to clarify the set-up thus far using a numerical example. Suppose that in the year 2018 an auction is held for year 2025 permits. The actual price that we will observe in 2025 is denoted P_{2025} and is a draw from $N(25, 50)$, where ρ equals 25 and Σ equals 50. This information is common knowledge. The actual price in 2025 will be $25 + \varepsilon$. No one knows this amount, however informed traders have private signals about ε . Therefore, traders using their signals will have different expectations regarding P_{2025} , such as (for example) 20, 22, 25, 27, and 29.

Now, suppose that a price p_k emerges from the auction for 2025 allowances. If both p_k and $\rho + \theta \bar{g}$ are the same, for instance if they both equal \$25.50, then the price arising from the auction is a sufficient statistic for all the available information. Auction participants cannot directly verify whether this is the case because they only observe the auction price p_k . Suppose, for instance, that $\rho + \theta \bar{g} = 25.50$ but $p_k = \$25.80$. This implies the auction price has not used all the available information.

What we will show subsequently is that both $p_k = \$25.50$ and $p_k = \$25.80$ are unbiased forecasts of P_{2025} given the order flow. The difference between the two prices reflects how much of their private signals agents choose to reveal through trading activity. By looking just at the auction price p_k there is no way to know whether it is a sufficient statistic. Sufficiency can only be inferred based on the information structure of the market. It will turn out to depend critically on Ω_0 , namely the initial correlation of the signals traders draw upon, which we interpret as a measure of the degree of scientific consensus on climate change.

Uninformed traders submit a quantity u_h at round h that has a normal distribution with mean zero and variance σ_u^2 . We assume uninformed traders' order u_h to be independent of all the other random variables. They have no private information and nothing can be learned about p_T from their actions during the trading process.

The i^{th} informed trader, given his private information and what he learns from past prices, submits a quantity q_{ih} at round h . Market makers observe the aggregate quantity y_h submitted by both informed and uninformed traders, update their beliefs regarding the initial signals using $g_{ih} = E[g_{i0} | y_1, \dots, y_h]$,⁶ and set the competitive price at period h using $p_{hT} = E[p_T | y_1, \dots, y_h]$, meaning that market makers earn zero profits on average, by assumption.

In this market, information gets fully aggregated if the market price contains information from all market traders so that each informed trader finds his own private information, g_{ih} , redundant. In other words, private information gets fully aggregated if a round h occurs in which

$$p_{hT} = E[p_T | y_1, \dots, y_h] = E(p_T | g_{10}, g_{20}, \dots, g_{M0}) = \rho + \theta \bar{g} \quad (2)$$

Therefore, if the auction price becomes equal to $\rho + \theta \bar{g}$, it becomes a perfect aggregator of information, and consequently an ideal estimate of p_T . In this case, the price contains all the available information and becomes fully revealing, without containing any error in predicting p_T .⁷ Note that traders do not necessarily know if the observed price p_{hT}

⁶ Market makers need to update their estimate regarding initial signals so that they can have a better estimate of the true price of a permit p_T . In particular, price, which is the market makers' updated estimate of p_T given their observed order flow, is related to the market makers' estimate of signals in the following way (FV96)

$$p_{hT} = \frac{\theta}{M} \sum_{i=1}^M \bar{g}_{ih}$$

⁷ In particular, we assume that the ideal estimate of p_T , $\rho + \theta \bar{g}$, is indeed equal to p_T .

satisfies (2) or not. They only observe the price in this market, not the underlying “true” price $\rho + \theta\bar{g}$. As a result, the magnitude by which the price at an auction, which is also an estimate of p_T , differs from $\rho + \theta\bar{g}$, shows the level of noise, and is of interest in this analysis. We define noise Σ_h as the following:

$$\Sigma_h = \text{Var}(\rho + \theta\bar{g} - p_{hT}) \quad (3)$$

Σ_h measures the magnitude of noise at round h in an auction system conveying information regarding the expected price of the permit. In other words, it measures informativeness at each round h in the climate prediction market. If Σ_h equals zero, that would mean that the price at round h has become exactly equal to $\rho + \theta\bar{g}$, meaning that all the information has been aggregated and incorporated into the price. In other words, the auction price is an unbiased estimate of the true value p_T , though it contains some noise and the magnitude of this noise shows the extent to which prices become informative in this market.⁸

After h rounds market makers observe (y_1, \dots, y_h) and learn about private signals, so at round h , the i^{th} informed player has an informational advantage relative to market makers that is equal to:

$$g_{ih} = g_{i0} - E[g_{i0} | y_1, \dots, y_h] \quad (4)$$

FV96 also define the following variance covariances in order to measure the remaining information:

$$\Lambda_h = \text{Var}(g_{i0} | y_1, \dots, y_h) = \text{Var}(g_{ih}) \quad (5)$$

$$\Omega_h = \text{Cov}(g_{j0}, g_{i0} | y_1, \dots, y_h) = \text{Cov}(g_{jh}, g_{ih}) \quad (6)$$

where i and j are two different traders.

4. The equilibrium

We have used the FV96 and Kyle (1985) auction model to characterize a permits trading system that would implement the Hsu (2011) exemption allowance proposal. Now we need to characterize the equilibrium that emerges and explore the properties of the resulting prices as they relate to the question of whether such a market would yield valid climate forecasts. Following Kyle (1985) and FV96, a Bayesian Nash equilibrium exists if there is a vector of strategies $(q_1, q_2, \dots, q_M, p)$ such that:

1. For each trading round $h = 1, 2, \dots, k$ and for every informed trader i , $i = 1, 2, \dots, M$, an alternative strategy $q_i' = (q_{i1}', q_{i2}', \dots, q_{ih}')^T$ would result in a lower or equal expected profit for firm i given this information set, the pricing rule, and strategies of the other informed traders:

$$E(\pi_h(q_1, q_2, \dots, q_i', \dots, q_M, p) | g_{i0}, q_{i-h}, y_{-h}) \leq E(\pi_h(q_1, q_2, \dots, q_i, \dots, q_M, p) | g_{i0}, q_{i-h}, y_{-h}) \quad (7)$$

i.e. The optimal strategy of trader i at period h should be best no matter what strategies he played in previous periods; and

2. The price at round $h = 1, 2, \dots, k$ becomes equal to the expected value of the permit given the observed order flow up to h ; and given the strategies of the informed traders.

$$p_{hT} = E[p_T | y_1, y_2, \dots, y_h] \text{ for } h = 1, 2, \dots, k \quad (8)$$

This is a market efficiency condition that makes price become an unbiased estimate of the future price and hence the future tax rate. This condition can be viewed as being a result of a Bertrand auction among at least two risk-neutral market makers who can observe only the

total quantities in the market. The Bertrand auction would result in a price which yields zero profits for market makers, indicating that price becomes equal to the permit's expected value.

FV96 used dynamic programming and backward induction to characterize a linear equilibrium in this setting. Proposition 1 restates their result in the context of the application developed herein, showing the existence of an equilibrium.

4.1. Proposition 1

There exists a recursive linear Markov equilibrium that satisfies conditions 1 and 2. The equilibrium occurs where all informed traders $i = 1, 2, \dots, M$ submit bids of the form $q_{ih} = \beta_h g_{ih}$ for all trading periods $h = 1, 2, \dots, k$, and prices become equal to $p_h = p_{h-1} + \lambda_h y_h$, where the parameters β_h and γ_h are defined as follows:

$$\beta_h = \frac{\eta_h - \lambda_h \Psi_h}{\lambda_h [1 + (1 + \phi_h(M-1))(1 - (\lambda_h \Psi_h / \theta))]}$$

$$\gamma_h = \frac{(1 - 2\mu_h \lambda_h)(1 - (\lambda_h \beta_h(M-1)/\theta))}{2\lambda_h(1 - \mu_h \lambda_h)}$$

$$\alpha_{h-1} = (\eta_h - \lambda_h \beta_h [1 + (M-1)\phi_h])\beta_h + \alpha_h \left[1 - \frac{\lambda_h \beta_h}{\theta} [1 + (M-1)\phi_h]\right]^2$$

$$\Psi_{h-1} = (\eta_h - \lambda_h \beta_h [1 + (M-1)\phi_h])\gamma_h - \lambda_h \gamma_h \beta_h + \beta_h \left(1 - \frac{\lambda_h \beta_h(M-1)}{\theta}\right) + \Psi_h \left(1 - \frac{\lambda_h \beta_h}{\theta} [1 + (M-1)\phi_h]\right) \left(1 - \frac{\lambda_h \beta_h(M-1)}{\theta} - \lambda_h \gamma_h\right)$$

$$\mu_{h-1} = -\lambda_h \gamma_h^2 + \gamma_h \left(1 - \frac{\lambda_h \beta_h(M-1)}{\theta}\right) + \mu_h \left(1 - \frac{\lambda_h \beta_h(M-1)}{\theta} - \lambda_h \gamma_h\right)^2$$

$$\delta_{h-1} = \delta_h + \frac{\alpha_h \lambda_h^2}{\theta^2} \sigma_u^2 + \alpha_h \frac{\beta_h^2 \lambda_h^2}{\theta^2} (M-1) \text{Var}(g_{jh-1} | g_{i0}, y_1, \dots, y_{h-1}) + \alpha_h \frac{\beta_h^2 \lambda_h^2}{\theta^2} [(M-2)(M-1)] \text{Cov}(g_{jh-1}, g_{m-1} | g_{i0}, y_1, \dots, y_{h-1})$$

$$\text{Var}(g_{jh-1} | g_{i0}, y_1, \dots, y_{h-1}) = \frac{\Lambda_{h-1}^2 - \Omega_{h-1}^2}{\Lambda_{h-1}}$$

$$\text{Cov}(g_{jh-1}, g_{rh-1} | g_{i0}, y_1, \dots, y_{h-1}) = \frac{\Omega_{h-1}(\Lambda_{h-1} - \Omega_{h-1})}{\Lambda_{h-1}}$$

$$\phi_h = \frac{\Omega_{h-1}}{\Lambda_{h-1}}$$

$$\eta_h = \frac{\theta}{M} [1 + (M-1)\phi_h]$$

$$\Sigma_h = \left(1 - \frac{M\lambda_h \beta_h}{\theta}\right) \Sigma_{h-1}$$

$$\Lambda_h = \Lambda_{h-1} - \frac{M\lambda_h \beta_h}{\theta^2} \Sigma_{h-1}$$

$$\Omega_h = \Omega_{h-1} - \frac{M\lambda_h \beta_h}{\theta^2} \Sigma_{h-1}$$

where $\alpha_h = \Psi_h = \mu_h = \delta_h = 0$.

A second order condition ensuring that the informed traders' utility is maximized also holds:

$$\lambda_h(1 - \mu_h \lambda_h) > 0$$

Proof: see proof of proposition 1, FV96 Appendix.

⁸ FV96 assumed ρ to be equal to zero so that they defined Σ_h as the variance of $(\theta\bar{g} - p_{hT})$.

The stated conditions are the necessary and sufficient conditions for the Markov equilibrium to hold. The stated recursions were solved numerically in FV96 using a backward induction algorithm and the equilibrium parameters were calculated for different correlation structures between the initial signals.

Before proceeding to the next section, it is worth clarifying the role of uninformed traders in this model. As stated in the above equilibrium, prices are a linear function of the order flows and become equal to:

$$p_h = p_{h-1} + \lambda_h y_h \quad (9)$$

In period $h = 1$ we will have:

$$p_1 = p_0 + \lambda_1 y_1 = \rho + \lambda_1 \left(\sum_{i=1}^M \beta_i g_{i1} + u_1 \right) \quad (10)$$

$$\implies p_1 = \rho + \lambda_1 \left(\beta_1 M \frac{\sum_{i=1}^M g_{i1}}{M} + u_1 \right) \quad (11)$$

$$\implies p_1 = \rho + \lambda_1 (\theta \bar{g} + u_1) \quad (12)$$

Therefore, in the equilibrium, as stated before and as can be seen from Eq. (12), price is a noisy signal of $\theta \bar{g}$ – a sufficient statistic for all the available information – and the source of the noise is coming from the uninformed traders' market order, u_1 . More specifically, uninformed traders in this model act as a camouflage for informed traders' information and help informed traders hide their information from market makers when trading with them. If we remove uninformed traders from the model, the price becomes fully revealing as it incorporates and reveals all the available information. In other words, removing uninformed traders from the model implies that prices in financial markets always reveal all the available information perfectly. This would lead to a paradox pointed out by Grossman and Stiglitz, 1980. When the equilibrium price is always a perfect aggregator of information, no traders would have any incentive to collect costly information as they know they cannot earn a return on their information gathering. Therefore, no one would gather information on which prices are based.

5. Properties of the market outcome

In this section, we present three properties of the market outcome that emerge from the analysis of the Bayesian Nash trading equilibrium.

1. The auction price $p_{kT} = E[p_T | y_1, y_2, \dots, y_k]$, is an unbiased estimate of the true future price p_T and hence of s_T . The unbiasedness of the price holds because:

$$E(p_{kT}) = E[p_T | y_1, y_2, \dots, y_k] = E[p_T] \quad (13)$$

Therefore, the unbiasedness condition $E(p_{kT} - p_T) = 0$ holds. In any specific outcome, prices are not necessarily equal to the true value, because they contain some noise, but the average of the existing noise is equal to zero.

2. The stronger the scientific "consensus" about global warming, the more accurate will be the forecasts implied by the auction price signals. More formally, the higher the correlation of signals received by traders, and hence the more correlated their beliefs, the closer the auction price gets to $\rho + \theta \bar{g}$, meaning it is a sufficient statistic, or one that uses all available information efficiently.

This was demonstrated by FV96 through numerical simulations of the equilibrium conditions. FV96 considered the case of four auctions ($k = 4$) and three informed traders ($M = 3$) in order to solve their model. They fixed the total available information and examined

four different information structures among traders. Specifically, they considered the following cases for the initial correlation among signals: very high positive (0.9999), which corresponds to identical information, low positive (0.1818), zero, and low negative (-0.2857). The evolution of Σ , which is the variance of the noise in this market, was computed for the four different cases. Fig. 1 is drawn using data from FV96, leaving out the negative correlation case since it does not apply in this model. It shows how Σ changes, or in other words how informative prices become, after four auctions under different signal structures.

As Fig. 1 shows, the higher the initial correlation, the lower the terminal variance and therefore the higher the price informativeness. The amount of reduction in sigma is very dramatic when traders have identical information, indicating that most of the information gets released after running only four rounds. But even when there is no initial agreement ($corr = 0$) the variance falls by more than half after four rounds. Results from Fig. 1 imply that traders with almost identical information in the prediction market would face more competition and trade more intensely in early rounds, causing more information to get released. However, traders with more heterogeneous information face less competition because part of their information is unique and gives them some monopoly power. As a result, they have less incentive to trade aggressively in early auctions, causing less information to get released over four auctions.

3. An auction price p_{kT} is a sufficient statistic (Eq. (1) holds) for p_T only when informed traders have almost identical private information regarding the state variable.

It was shown before that $\theta \bar{g}$ is a sufficient statistic for the available information to predict p_T (Eq. (1)). In other words, once $\theta \bar{g}$ is known, no further information can be gained about p_T by knowing any individual signals g_1, g_2, \dots, g_M . Therefore, an auction price p_{kT} becomes a sufficient statistic for p_{kT} when it becomes equal to $\rho + \theta \bar{g}$. As was shown in Fig. 1, p_{kT} becomes equal to $\rho + \theta \bar{g}$ or equivalently the level of noise becomes equal to zero even after only 4 trading rounds when informed traders had almost identical information in the market. But simulations in FV96 show that, with weaker correlation signals, Σ converges to zero but it remains positive even after 800 auction rounds. This means that repeated trading rounds may not compensate for information heterogeneity if prior beliefs are sufficiently uncorrelated.

In summary, the climate prediction market we have described for a future period T yields an unbiased estimate of the true future price, and hence of the climate state. Also, the level of consensus about climate science strongly influences the effectiveness of the market. The higher the correlation of information signals prior to trading, the more information is aggregated by the price forecast during trading. Adding more rounds of trading helps increase the accuracy of prices even when signals are uncorrelated. Put another way, the argument that a price emerging from an allowance auction as described herein systematically ignores important information about the future climate relies on the assumption that there is no scientific consensus around climate change. Belief that such a consensus exists implies that the price emerging from the market would use all available information and therefore would be the most informative climate forecast possible.

6. Extension to risk averse firms

The previous results rely on the assumption that traders are risk-neutral. Here we extend the framework by considering M informed agents who possess different private information and who are risk averse, using the following CARA (constant absolute risk aversion) utility functions with "A" as a risk aversion coefficient:

$$U(W_{N+1}) = -e(-AW_{N+1})$$




Fig. 1. Evolution of sigma with 4 periods, where zero indicates convergence on a full-information outcome. (Data source: FV96).

where W_{N+1} represents terminal wealth. The following proposition establishes the existence of a trading equilibrium in this case.

6.1. Proposition 2

There is an equilibrium where informed traders submit the optimal order q_i , that is a linear function of initial signals and is given by:

$$q_i = \beta g_{i0}$$

and the price is set according to the following linear rule:

$$p = \chi \left(\sum_{i=1}^M q_i + u \right) = \lambda y$$

where β and λ are given by:

$$\beta = \frac{\eta}{2\lambda + \phi\lambda(M-1) + A\phi}$$

$$\lambda = \frac{\frac{BM}{\theta} \Sigma_0}{\left(\frac{BM}{\theta}\right)^2 \Sigma_0 + \sigma_u^2}$$

The proof is in the [Appendix](#).

To check how informative prices become after only one auction round we need to find Σ_1 which is equal to $\text{Var}(p_T|y)$ by definition.

$$\Sigma_1 = \text{Var}(p_T|y) = \text{Var}(\bar{g}\theta|y) = \Sigma_0 (1 - \text{corr}_{\bar{g},y}^2) =$$

$$\Sigma_0 \left(1 - \frac{\text{Cov}(\bar{g}\theta, y)^2}{\text{Var}(\bar{g}\theta) \text{Var}(y)} \right) = \Sigma_0 \left(1 - \frac{\text{Cov}\left(\bar{g}\theta, \frac{\beta M}{\theta}(\bar{g}\theta) + u\right)^2}{\Sigma_0 \text{Var}\left(\frac{\beta M}{\theta}(\bar{g}\theta) + u\right)} \right)$$

$$= \Sigma_0 \left(1 - \frac{\left(\frac{\beta M}{\theta}\right)^2 \Sigma_0^2}{\Sigma_0 \left[\left(\frac{\beta M}{\theta}\right)^2 \Sigma_0 + \sigma_u^2 \right]} \right)$$

Therefore,

$$\Sigma_1 = \Sigma_0 \left(\frac{\sigma_u^2}{\left(\frac{\beta M}{\theta}\right)^2 \Sigma_0 + \sigma_u^2} \right)$$

As was proved, the i^{th} informed risk averse trader would trade $q_i = \beta g_{i0}$ in the first auction where $\beta = \frac{\eta}{2\lambda + \phi\lambda(M-1) + A\phi}$. Since $A\phi$ in the denominator is a positive term, comparing the optimal demand in the case of risk averse traders with the case of risk-neutral traders indicates that risk averse traders would trade smaller amounts, causing prices to be less informative. As it was shown, Σ_1 is given by:

$$\Sigma_1 = \Sigma_0 \left(\frac{\sigma_u^2}{\left(\frac{\beta M}{\theta}\right)^2 \Sigma_0 + \sigma_u^2} \right) = \Sigma_0 \left[\frac{\sigma_u^2}{\left(\frac{\eta}{2\lambda + \phi\lambda(M-1) + A\phi} \right)^2 \left(\frac{M}{\theta} \right)^2 \Sigma_0 + \sigma_u^2} \right]$$

Therefore, in the above equation $\frac{\partial \Sigma_1}{\partial A}$ is positive and is given by:

$$\frac{\partial \Sigma_1}{\partial A} = \frac{\partial \Sigma_1}{\partial \beta} \cdot \frac{\partial \beta}{\partial A} = \frac{-2 \left(\frac{M}{\theta} \right)^2 \Sigma_0 \beta (\Sigma_0 \sigma_u^2)}{\left[\left(\frac{\beta M}{\theta} \right)^2 \Sigma_0 + \sigma_u^2 \right]^2} \cdot \left(\frac{-\phi \eta}{[2\lambda + \phi\lambda(M-1) + A\phi]^2} \right)$$

$$= \frac{2 \left(\frac{M}{\theta} \right)^2 \Sigma_0^2 \sigma_u^2 \phi \eta \beta}{\left[\left(\frac{\beta M}{\theta} \right)^2 \Sigma_0 + \sigma_u^2 \right]^2 [2\lambda + \phi\lambda(M-1) + A\phi]^2}$$

As traders become more risk averse, they will trade smaller quantities and as a result, the permit price would become informative at a slower rate. In other words, we expect to see a lower rate of reduction in Σ when adding risk aversion to the model.

7. Conclusion

In this paper we developed the idea of integrating a state-contingent emission pricing rule with tradable permit markets as a way of obtaining an accurate forecast of the future climate state. We modeled trading of permits that exempt emitters from paying a temperature-indexed tax among traders with different private signals and investigated price formation in this market. With risk neutral traders, we used the FV96 framework to show how information gets efficiently aggregated and incorporated into prices, making them unbiased forecasts of the future climate outcomes. Also, we demonstrated that the initial structure of private information plays an important role in the level of competition among traders and consequently the extent to which prices become informative. In particular, the higher the level of scientific consensus (meaning the higher the correlation of private signals received by informed traders), the more competition traders face, causing more information to get released. With almost identical information, it takes very few auction rounds for all the information to get incorporated into the market price. On the other hand, traders with more heterogeneous information face less competition as they have some monopoly power. They trade less intensely, causing information to get released more slowly over time. We showed that, in this case, adding multiple rounds of bidding would help increase the price informativeness. We also extended the existing model to include risk averse traders. In this case, they trade smaller quantities and as a result, the permit price would become informative at a slower rate.

Note that we are assuming there is a minimal feedback between policy and the state variable. It is conceivable that if a steep warming path led to a path of steeply-rising emission taxes, and the response to the policy caused the warming path to drop, there could emerge multiple equilibrium paths. This is not likely to happen, however, due to the unresponsiveness of the climate even to major policy initiatives. Wigley (1998) and Lomborg (2016) modeled the climatic effects of, respectively, the international commitments under the Kyoto and Paris climate accords and showed that they barely budged the temperature path through 2100. Consequently, any individual nation implementing a state-contingent emission price would likely not expect its policy responses to have a large effect on the state variable.

These results are important because they provide a feasible mechanism for addressing two serious impediments to climate policy formation: uncertainty about the optimal future emission price given an available information set, and uncertainty about whether the information set itself is biased or incomplete. The combined state-contingent tax/exemption permit auction would yield a forward-looking price path that correlates to the unobservable inter-temporal marginal damages path, without a regulator needing to know the current form of the underlying state function connecting the climate to emissions. The permits auction would yield unbiased predictions of the future climate state. And it would make more efficient use of all available information, the more correlated are private information signals, in other words the greater the strength of the scientific consensus on climate change. If the consensus is weak, information would still aggregate, only more slowly, but that would correspond to the case in which the basis for any policy stance is likewise weakened. Finally, risk-aversion among traders would slow the rate of information aggregation but would not prevent it.

Acknowledgment

The authors thank the Institute for New Economic Thinking for financial support.

Declarations of interest

None.

Appendix A

Proof of Eq. (1):

Let X and Y be two random variables with normal distributions. The conditional expectation of X given $Y = y$ is the following formula:

$$E[X|Y = y] = m_x + \Sigma_{XY} \Sigma_{YY}^{-1} (Y - m_Y)$$

where $m_x = E[X]$, $m_Y = E[Y]$, and Σ is the covariance matrix. Given the above formula and the normality assumptions we will have:

$$E(p_T | g_{10}, g_{20}, \dots, g_{M0}) = m_{p_T} + \Delta'_0 \Psi_0^{-1} (g - m_g)'$$

p_T has a normal distribution with mean ρ_0 and variance Σ_0 . In the above equation, g is the vector of signals with zero means. The diagonal elements of Ψ_0^{-1} are all Λ_0 (initial variance of signals) and the off diagonal elements of Ψ_0^{-1} are all Ω_0 (initial covariance between any two signals). Δ_0 is the covariance vector between signals and the true price of a permit with c_0 as its element, meaning that $\Delta_0' = [c_0, c_0, \dots, c_0]$. Therefore the entries of the matrix $\Delta_0' \Psi_0^{-1}$ are all the same and equal to $c_0[\Lambda_0 + (M - 1)\Omega_0]$. We show the term $c_0[\Lambda_0 + (M - 1)\Omega_0]$ by ς . Thus,

$$E(p_T | g_{10}, g_{20}, \dots, g_{M0}) = \rho_0 + \varsigma \sum g_{i0} = \rho_0 + M\varsigma \left(\frac{\sum g_{i0}}{M} \right) = \rho_0 + \theta \bar{g}$$

where $\theta = M\varsigma$, which is a set of weights implied by conditioning on the signals.

Proof of Proposition 2

The i^{th} informed trader maximizes his expected utility given his initial signal and given the assumption that other traders follow their optimal strategies. Therefore, at round $h = 0$, he solves the following problem:

$$\begin{aligned} \text{Max}_{q_i} \quad & E[-e(-AW_1) | g_{i0}] = -E[e(-AW_0 - Aq_i(p_T - p)) | g_{i0}] \\ & = -e^{-AW_0} E \left[e \left(p_T - \lambda q_i - \lambda \sum_{i \neq j} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \end{aligned} \quad (14)$$

given that $p = \lambda(\sum_{i=1}^M q_i + u)$ and $\sum_{i \neq j} q_j = \beta g_{j0}$

We assume that all informed traders have the same initial wealth W_0 . Using the Linear-Normal-Exponential model properties, we will have the following:

$$\begin{aligned} & E \left[\exp \left(p_T - \lambda q_i - \lambda \sum_{i \neq j} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \\ & = \exp \left\{ E \left[\left(p_T - \lambda q_i - \lambda \sum_{i \neq j} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \right. \\ & \quad \left. + \frac{1}{2} \text{Var} \left[\left(p_T - \lambda q_i - \lambda \sum_{i \neq j} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \right\} \end{aligned} \quad (15)$$

Therefore, maximizing Eq. (14) is equivalent to maximizing:

$$\begin{aligned} & E \left[\left(p_T - \lambda q_i - \lambda \sum_{i \neq j} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \\ & + \frac{1}{2} \text{Var} \left[\left(p_T - \lambda q_i - \lambda \sum_{i \neq j} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \end{aligned} \quad (16)$$

since the expression in Eq. (16) is a monotone increasing transformation of the expression in (15).

From Eq. (16):

$$\begin{aligned} E \left[\left(p_T - \lambda q_i - \lambda \sum_{j \neq i} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \\ = -Aq_i \left\{ E[p_T | g_{i0}] - \lambda q_i - \lambda \sum_{j \neq i} E[\beta g_{j0} | g_{i0}] \right\} \end{aligned} \quad (17)$$

Multivariate normality implies that both $E[p_T | g_{i0}]$ and $E[g_{j0} | g_{i0}]$ are linear in g_{i0} , therefore Eq. (17) becomes equal to:

$$-Aq_i [\eta g_{i0} - \lambda q_i - \lambda \beta (M-1) \varphi g_{i0}]$$

where

$$\begin{aligned} \eta &= \frac{\text{Cov}(p_T, g_{i0})}{\text{Var}(g_{i0})} = \frac{\text{Cov}(\theta \bar{g}, g_{i0})}{\text{Var}(g_{i0})} = \frac{\theta \text{Cov}(\frac{\sum g_{i0}}{M}, g_{i0})}{\text{Var}(g_{i0})} = \frac{\theta}{M} [\Lambda_0 + (M-1)\Omega_0] \\ &= \frac{\theta}{M} \left[1 + \frac{(M-1)\Omega_0}{\Lambda_0} \right] \end{aligned}$$

and

$$\varphi = \frac{\text{Cov}(g_{j0}, g_{i0})}{\text{Var}(g_{i0})} = \frac{\Omega_0}{\Lambda_0}$$

and

$$\begin{aligned} \text{Var} \left[\left(p_T - \lambda q_i - \lambda \sum_{j \neq i} q_j - \lambda u \right) (-Aq_i) | g_{i0} \right] \\ = A^2 q_i^2 \left[\text{Var}(p_T | g_{i0}) + \lambda^2 \beta^2 \text{Var} \left(\sum_{j \neq i} g_{j0} | g_{i0} \right) + \lambda^2 \sigma_u^2 \right] = A^2 q_i^2 \vartheta \end{aligned}$$

where

$$\vartheta = [\text{Var}(p_T | g_{i0}) + \lambda^2 \beta^2 \text{Var}(\sum_{j \neq i} g_{j0} | g_{i0}) + \lambda^2 \sigma_u^2].$$

Therefore, the i^{th} informed player's problem is to maximize

$$-Aq_i [\eta g_{i0} - \lambda q_i - \lambda \beta (M-1) \varphi g_{i0}] + \frac{1}{2} A^2 q_i^2 \vartheta$$

with respect to q_i .

The first-order condition is:

$$-A[\eta g_{i0} - \lambda q_i - \lambda \beta (M-1) \varphi g_{i0}] + (-\lambda)(-Aq_i) + A^2 \vartheta q_i = 0$$

$$q_i (2A\lambda + A^2 \vartheta) = A\eta g_{i0} - A\lambda \beta (M-1) \varphi g_{i0}$$

$$q_i = \frac{\eta - \lambda \beta (M-1) \varphi}{2\lambda + A\vartheta} g_{i0}$$

Therefore q_i is a linear function of the initial signal, and β is given by:

$$\beta = \frac{n}{2\lambda + \phi\lambda(M-1) + A\vartheta}$$

Now with respect to the price:

$$p = E[p_T | y] = E \left[p_T \left| \sum_{i=1}^M q_i + u \right. \right] = E \left[p_T \left| \sum_{i=1}^M \beta g_{i0} + u \right. \right] = \lambda y$$

As both p_T and $y (= \sum_{i=1}^M \beta g_{i0} + u)$ are normally distributed, $E[p_T | y]$ becomes linear in y and λ is given by:

$$\lambda = \frac{\text{Cov}(p_T, y)}{\text{Var}(y)}$$

where

$$y = \sum_{i=1}^M \beta g_{i0} + u = \beta M \left(\frac{\sum_{i=1}^M g_{i0}}{M} \right) + u = \frac{\beta M}{\theta} (\bar{g}\theta) + u$$

⇒

$$\begin{aligned} \lambda &= \frac{\text{Cov} \left(p_T \frac{\beta M}{\theta} (\bar{g}\theta) + u \right)}{\text{Var} \left(\frac{\beta M}{\theta} (\bar{g}\theta) + u \right)} = \frac{\text{Cov} \left(\bar{g}\theta \frac{\beta M}{\theta} (\bar{g}\theta) + u \right)}{\text{Var} \left(\frac{\beta M}{\theta} (\bar{g}\theta) + u \right)} \\ &= \frac{\frac{\beta M}{\theta} \text{Var}(\bar{g}\theta)}{\left(\frac{\beta M}{\theta} \right)^2 \text{Var}(\bar{g}\theta) + \text{Var}(u)} = \frac{\frac{\beta M}{\theta} \Sigma_0}{(\beta M / \theta)^2 \Sigma_0 + \sigma_u^2} \end{aligned}$$

This completes the proof.

References

- Atakan, Alp E., Ekmekci, Mehmet, 2014. Auctions, actions and the failure of information aggregation. *Am. Econ. Rev.* 104 (7), 2014–2048.
- Berg, Joyce E., Nelson, Forrest D., Rietz, Thomas A., 2008. Prediction market accuracy in the long run. *Int. J. Forecast.* 24 (2), 285–300.
- Berry, Brian J.L., Bihari, Jayshree, Elliott, Euel, 2016. The limits of knowledge and the climate change debate. *Cato J.* 36 (3), 589–610.
- Winton Capital, 2017. Winton and LSE Researchers Propose UK-based Climate Prediction Market. <https://www.winton.com/news/2017/climate-prediction-market> (website accessed November 14, 2017).
- Cederlöf, Magnus, Bengtsson, Lennart, Hodges, Kevin, 2016. Assessing atmospheric temperature data sets for climate studies. *Tellus A* 68, 31503. <https://doi.org/10.3402/tellusa.v68.31503>.
- Fama, Eugene F., 1970. Efficient capital markets: a review of theory and empirical work. *J. Financ.* 25 (2), 383–417.
- Foster, F. Douglas, Viswanathan, S., 1996. Strategic trading when agents forecast the forecasts of others. *J. Financ.* 1437–1478.
- Graefe, Andreas, 2016. Political markets. In: Arzheimer, Kai, et al. (Eds.), *The Sage Handbook of Electoral Behaviour*, p. 2016.
- Graefe, Andreas, 2017. Prediction market performance in the 2016 U.S. presidential election. *Foresight Int. J. Appl. Forecasting* 45, 38–42.
- Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. *Am. Econ. Rev.* 70 (3), 393–408.
- Hayek, F.A., 1945. The use of knowledge in society. *Am. Econ. Rev.* 519–530.
- Hsu, Shi-Ling, 2011. A prediction market for climate outcomes. *Univ. Colorado Law Rev.* 83, 179–256.
- InterAcademy Council (IAC), 2010. "Climate Change Assessments: Review of the Processes and Procedures of the IPCC." Netherlands: IAC, October 2010. <http://reviewipcc.interacademycouncil.net/report.html>.
- Johnston, Jason Scott, 2012. In: Johnston, Jason Scott (Ed.), *The Costs of Cartelization: The IPCC Process and the Crisis of Credibility in Climate Science. Institutions and Incentives in Regulatory Science*, Lexington.
- Kelly, David L., Kolstad, Charles D., 1999. Bayesian learning, growth, and pollution. *J. Econ. Dyn. Control.* 23 (4), 491–518.
- Kyle, Albert S., 1985. Continuous auctions and insider trading. *Econometrica* 1315–1335.
- Leach, Andrew J., 2007. The climate change learning curve. *J. Econ. Dyn. Control.* 31 (5), 1728–1752.
- Lemoine, Derek M., Traeger, Christian P., 2012. Tipping points and ambiguities in the economics of climate change. NBER Working Paper 18230, July 2012.
- Lemoine, Derek M., Traeger, Christian P., 2016. Economics of tipping the climate dominoes. *Nat. Clim. Chang.* 6:514–519. <https://doi.org/10.1038/nclimate2902>.
- Lomborg, Bjorn, 2016. Impact of current climate proposals. *Global Pol.* 7 (1), 109–118.
- McKittrick, Ross, 2011. A simple state-contingent pricing rule for complex intertemporal externalities. *Energy Econ.* 33 (1), 111–120.
- Nordhaus, William D., 1991. To slow or not to slow: the economics of the greenhouse effect. *Econ. J.* 920–937.
- Nordhaus, William D., 1993a. Optimal greenhouse-gas reductions and tax policy in the "DICE" model. *Am. Econ. Rev.* 313–317.
- Nordhaus, William D., 1993b. Rolling the 'DICE': an optimal transition path for controlling greenhouse gases. *Resour. Energy Econ.* 15 (1), 27–50.
- Roe, Gerard H., Baker, Marcia B., 2007. Why is climate sensitivity so unpredictable? *Science* 318 (5850), 629–632 (October 26, 2007).
- Roe, Gerard H., Bauman, Yoram, 2013. Climate sensitivity: should the climate tail wag the policy dog? *Clim. Chang.* 117:647–662. <https://doi.org/10.1007/s10584-012-0582-6>.

- Segol, Matthieu, 2012. Aggregation of information and beliefs on prediction markets with non-bayesian traders. No. Dumas-00809694.
- Tol, Richard S.J., 2014. Quantifying the consensus on anthropogenic global warming in the literature: a re-analysis. *Energy Policy* 73, 701–705 (October 2014).
- US Interagency Working Group on Social Cost of Carbon, 2013. Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866. United States Government.
- Weitzman, M., 2009. On modeling and interpreting the economics of catastrophic climate change. *Rev. Econ. Stat.* 91 (1), 1–19.
- Wigley, T.M.L., 1998. The Kyoto protocol: CO₂, CH₄ and climate implications. *Geophys. Res. Lett.* 25 (13), 2285–2288.
- Wolfers, Justin, Zitzewitz, Eric, 2004. Prediction markets. National Bureau of Economic Research Working Paper No. w10504.